Journal of Digital Imaging

Interactive Image Enhancement of CR and DR Images

Maikael A. Thomas, MSEE,[1] Alan H. Rowberg, MD,[2] Steve G. Langer, PhD,[3] Yongmin Kim, PhD[1]

There is continual pressure on the radiology department to increase its productivity. Two important links to productivity in the computed/digital radiography (CR/DR) workflow chain are the postprocessing step by technologists and the primary diagnosis step by radiologists, who may apply additional image enhancements to aid them in diagnosis. With the large matrix size of CR and DR images and the computational complexity of these algorithms, it has been challenging to provide interactive image enhancement, particularly on full-resolution images. We have used a new programmable processor as the main computing engine of enhancement algorithms for CR or DR images. We have mapped these algorithms to the processor, maximally utilizing its architecture. On a 12-bit 2688 × 2688 image, we have achieved the execution time of 465 ms for adaptive unsharp masking, window/level, image rotate, and lookup table operations using a single processor, which represents at least an order of magnitude improvement compared to the response time of current systems. This kind of performance facilitates rapid computation with preset parameter values and/or enables truly interactive QA processing on radiographs by technologists. The fast response time of these algorithms would be especially useful in a real-time radiology setting, where the radiologist's waiting time in performing image enhancements before making diagnosis can be greatly reduced. We believe that the use of these processors for fast CR/DR image computing coupled with the seamless flow of images and patient data will enable the radiology department to achieve higher productivity.

KEY WORDS: Digital radiography, computed radiography, real-time radiology, high-performance computing, workflow, image enhancement, QA processing

DESPITE BEING the oldest of medical imaging modalities, traditional projection radiography accounts for roughly 70% of all radiological examinations performed.[1] In the 1990s, the use of computed radiography (CR) became popular. Most CR systems use photostimulable phosphor imaging plates housed inside of cassettes, which must be taken to a separate CR reader to be scanned and form the digital image.[2] With advances in high-resolution flat panel detector technology, digital radiography (DR) systems have recently been commercially introduced. DR images are acquired by a flat panel detector and then immediately digitized. The main advantage of CR and DR in a picture archiving and communications system (PACS) environment over traditional film-based systems is the separation of image capture, processing, storage, and display, allowing for the individual optimization of each of these steps.[3]

With decreasing reimbursement rates and increasing competition among diagnostic imaging providers, many radiology departments are faced with a challenging situation to reduce staff and operational costs while maintaining or increasing the quantity and quality of imaging services. [4] Thus, many institutions are trying to increase the productivity of the radiology department. The workflow of a film-

[1]*From the Image Computing Systems Laboratory, Departments of Electrical Engineering and Bioengineering, University of Washington, Seattle, WA 98195, USA.*

[2]*From the Department of Radiology, University of Washington, Seattle, WA 98195, USA.*

[3]*From the Department of Radiology, Mayo Clinic, Rochester, MN 55905, USA.*

Correspondence to: Yongmin Kim, PhDtel: 206-685-2271; fax: 206-221-6837; e-mail: ykim@u.washington.edu

Online publication 4 June 2004

doi: 10.1007/s10278-004-1000-z

based department is quite inefficient compared with that of a filmless and paperless department. Siegel and Reiner [5] identified that 59 steps and 11 staff members are necessary for the ordering, acquisition, reporting, and review of inpatient chest radiographs using film-based radiography. They redesigned the workflow to utilize CR with the PACS integrated with the hospital/radiology information system (HIS/RIS) to facilitate the automated electronic transport of images and patient medical records. They were able to eliminate 50 steps from the film-based workflow, requiring only four staff members for the entire process. We illustrate the remaining nine steps in Figure 1, which applies to both CR and DR. In Step 5, the technologist executes enhancement (or postprocessing) algorithms to prepare the images for the upcoming diagnosis and consultation process. In Step 7, the radiologist may apply further image enhancements to aid primary diagnosis. The more time these steps take, the lower the throughput in the CR/DR imaging workflow. Further improvement in productivity could be achieved by providing radiology services in a real-time fashion, ie, delivery of the radiologist's interpretation almost immediately after the examination. Even though real-time radiology would increase the demand on radiologists, it would better serve the patient, especially in urgent situations (eg, intensive care and emergency room), by significantly decreasing the time to diagnosis.[6]

The image enhancement and manipulation algorithms typically performed on CR and DR images are computationally expensive. To run all these algorithms on CR/DR images has been a computational challenge. For example, the AGFA DR-Thorax system takes approximately 30 s for all image processing of a 2688×2688 image,[7] while the Kodak DirectView DR 5100 has a cycle time of up to 38 s for a 2560×3072 image.[8] These systems take from 3 to 10 s to generate a preprocessed image from the time of X-ray exposure. Considering other components included in the cycle time, it is safe to assume that the postprocessing algorithms take on the order of 10 s on these systems.

Multiple input parameters to image enhancement algorithms are possible, which may require tuning based on a variety of factors, such as X-ray dose, body part, patient size, viewer preference, or a specific pathology of interest. Kodak's enhanced visualization processing (EVP), [9]Fuji's dynamic range control (DRC),[10] and AGFA'S multiscale image contrast amplification (MUSICA)[11] are typical CR/DR enhancement algorithms, each with several adjustable input parameters. These algorithms tackle the problem of the diagnostically relevant dynamic range of a radiograph being too large to be displayed with good contrast. EVP and DRC each have three selectable parameters, while MUSICA has four. These parameters control different aspects in image enhancement, such as latitude, contrast, and edge enhancement. Due to the numerous possible combinations of these parameters, many institutions use presets in QA processing, which can simplify and standardize the image enhancement process.

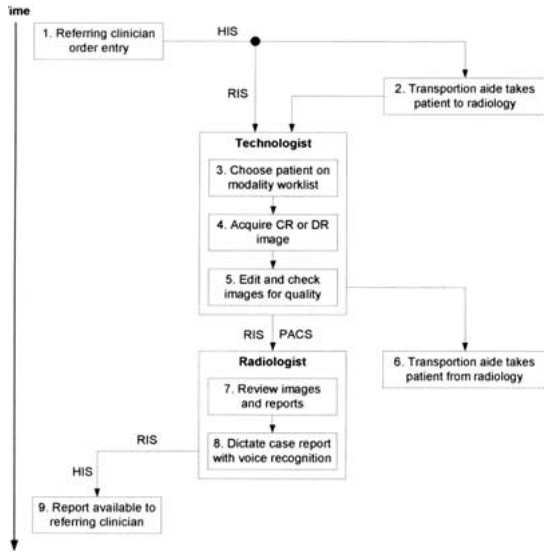

Fig 1. Nine steps of the redesigned CR or DR workflow for inpatient chest radiographs.

Our goal is to provide a computing solution that quickly executes image enhancement algorithms on CR and DR images for the technologist and radiologist. Meeting this objective will facilitate the fast computation of presets and/or allow truly interactive adjustment of enhancement parameters. A fast response time of these algorithms can reduce the burden on radiologists in a real-time radiology setting by decreasing the waiting time in performing image

enhancements while making diagnosis. For CR and DR computing, we have used a new generation of high-performance processors, developed for consumer products in imaging and video applications, known as mediaprocessors.[12] To demonstrate the feasibility of utilizing a mediaprocessor as a computing engine for CR/DR image processing algorithms, we have mapped representative postprocessing algorithms to a single MAP processor from Hitachi/Equator Technologies.[13]

METHODS

Image Enhancement Algorithms

Our computing module includes commonly found enhancement algorithms that optimize the radiograph for output display.[3] We have implemented two contrast enhancement algorithms: (1) window/level to rescale contrast and (2) adaptive unsharp masking to enhance mid- to high-frequency signals. As an additional contrast enhancement method, we have also included a general lookup table (LUT) function, applying a selected LUT to the image. LUTs are simple but useful in many image processing and enhancement applications. While grid suppression is typically a preprocessing algorithm, we have included it as a postprocessing algorithm to allow the technologist to interactively view the image with or without the grid. Also, we have included image rotation and flip functions. Figure 2 shows these algorithms as they apply to technologists and radiologists. Technologists check and edit radiographs for quality to reduce the amount of time radiologists must spend performing further manipulations. Radiologists can then focus more of their efforts on diagnosis, performing only a minimum amount of enhancement.

A linear grid is commonly used in projection radiography to attenuate scattered X-ray photons, thus improving the contrast and signal-to-noise ratio of images. When using a stationary grid, however, its lead strips cause a periodic shadow of grid lines to appear on the image. We determine the existence and exact frequency of the grid lines as described by Barski and Wang.[14] If a grid is not detected (eg, examinations where no grid or a moving grid is used), the grid suppression process is skipped and the image passes to the subsequent enhancement step. Otherwise, we commence grid suppression by creating a Gaussian bandpass filter based on the exact grid frequency that was determined. Via the one-dimensional (1D) Fast Fourier Transform (FFT), we convert each row of the image to the frequency domain. We filter the 1D FFT result with the bandpass filter, effectively removing almost all image information except grid-related data. The grid suppression algorithm then computes the 1D inverse FFT on each of the rows, producing an image containing only grid-related information. Finally, we subtract the grid image from the original image, creating the grid-suppressed image. The grid suppression process has little effect on the noise content of the image, except the removal of noise contained at and around the grid frequency. It is worthwhile to note that the grid frequency should not be close to the sampling frequency of the CR reader or DR detector. Otherwise, the aliased grid frequency might mix with important image content and grid suppression could significantly sacrifice image quality.[14]

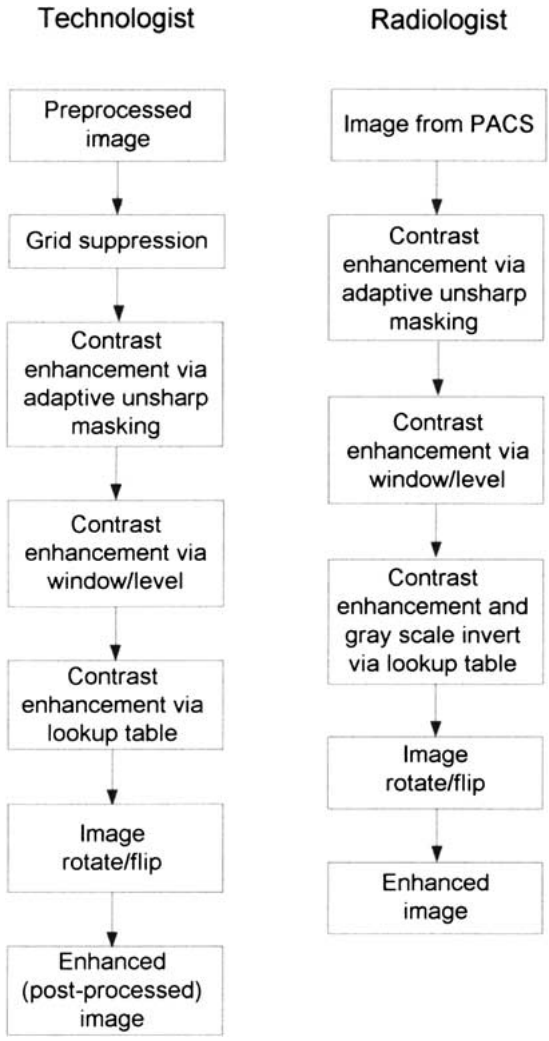

Fig 2. CR and DR enhancement algorithms for the technologist and radiologist.

We have used adaptive unsharp masking to increase image contrast by amplifying the mid- to high-frequency components of the image. We first convolve the input image with a boxcar filter. Using the boxcar filter is advantageous since the cutoff frequency can be made extremely low by using a large kernel size. Next, we obtain an image containing mid- and high-frequency components by subtracting the low-pass-filtered image from the original image. The resultant high-pass-filtered image is then scaled by the weighting factor and added to the original image. With global unsharp masking, the weighting factor is applied

equally across the image, which has the advantage of being simple. However, it is disadvantageous in that it amplifies noise in low X-ray density areas where the signal-to-noise ratio is lower. On the other hand, adaptive unsharp masking has the benefit of adaptively controlling the weighting factor for each pixel based on local image characteristics. This is useful because low-X-ray-density areas cannot be suppressed indiscriminately because an anatomical object of interest may exist, eg, lesions in thoracic or cervical spine. We calculate the local 3×3 gradient magnitude to emphasize both object contours and high-frequency isolated noise. Then, to discriminate objects from noise, we take the maximum gradient magnitude of the local 3×3 neighborhood. Spikelike noise tends to produce large gradient magnitude values compared with anatomical objects, whose pixel values change more gradually. Thus, if the maximum gradient magnitude value is above an adjustable threshold, the local region is likely to be noisy and we clip the weighting factor. Otherwise, the weighting factor value is set to be linearly proportional to the local gradient strength.

Mediaprocessor Architecture

Examples of commercially available mediaprocessors include the Hitachi/Equator Technologies MAP, the Texas Instruments TMS320C64, the TTI TriMedia, and the Intel Pentium 4.[15] Mediaprocessors were designed to excel in imaging and video applications where an operation is applied repetitiously across a large amount of data. Known as instruction-level parallelism, mediaprocessors have multiple independent functional units, each capable of executing an operation every clock cycle. Mediaprocessors also employ data-level parallelism, with which 32- or 64-bit operands (depending on mediaprocessor architecture) can be treated as a multiple of subwords (eg, 8, 16, or 32 bits). By performing partitioned arithmetic operations on these subwords, a 2×, 4×, or 8× speedup can be realized.[16] Of course, to achieve the maximum potential performance, all the functional units need to be utilized as much as possible.

In this study, we used the 450-MHz MAP mediaprocessor to compute CR/DR image enhancement algorithms. Many of our enhancement algorithms are based on key image processing algorithms, eg, convolution and FFT. We show how we have mapped the 1D real FFT and boxcar filtering to the MAP as an example of how to attain fast execution of the most computationally expensive portions of these algorithms.

Algorithm Mapping

A radix-2 N-point FFT requires $(2N\log_2 N - 4N + 4)$ real multiplications and $(3N\log_2 N - 2N + 2)$ real additions/subtractions.[17] Thus, a 4096-point 1D FFT performed on 2688 rows requires slightly less than 600 million operations. The MAP's *complex_multiply* instruction can issue two 16-bit complex multiplications in a single cycle, which amounts to eight multiplications and four additions/subtractions. Other instructions that are useful for FFT butterfly computations include *partitioned_add/subtract* to perform multiple 16-bit additions and subtractions and 64-bit *load* and *store*. Details on our FFT implementation can be found elsewhere.[18] Digital radiographs contain only real components, thus we can use the real FFT rather than the complex FFT to further reduce the amount of computation. Since the Fourier transform of a real signal is symmetric, the real FFT can be computed with an $N/2$-point complex FFT and a postprocessing step.[19] Furthermore, our real FFT generates only half of the frequency spectrum, which halves the amount of frequency-domain computations. Similarly, the real inverse FFT can be utilized.

A brute force approach to large-kernel convolution with large images is computationally very expensive, requiring $M^2 - 1$ additions for an $M \times M$ kernel size for each output point. Thus, convolving a 2688×2688 image with a 128×128 kernel would require about 235 billion arithmetic operations, which is impractical. Exploiting the separability property of the 2D boxcar filter reduces this number to $2(M - 1)$ additions. To maintain the filtered image size the same as the input image size, we pad the input image with boundary image pixels before we begin performing boxcar filtering. We begin by defining the initial row-wise pixel average over M pixels as

$$f_{\mathrm{ma}}(0,j) = \frac{1}{M} \sum_{k=-(M-1)/2}^{(M-1)/2} f_{\mathrm{pad-in}}(k,j)$$

where $f_{pad\text{-}in}$ is the padded input image and M is the odd-numbered kernel size. We compute the next point $f_{\mathrm{ma}}(1, j)$ by using the initial moving average $f_{\mathrm{ma}}(0, j)$ via

$$f_{\mathrm{ma}}(1,j) = f_{\mathrm{ma}}(0,j) + f_{\mathrm{pad-in}}\left(\frac{(M-1)}{2} + 1, j\right) - f_{\mathrm{pad-in}}\left(-\frac{(M-1)}{2}, j\right)$$

creating a new moving average. All subsequent intermediate output values are computed similarly by taking the previous moving average and calculating a new moving average according to Eq. (2). Thus, all intermediate outputs after the initial moving average require only one addition and one subtraction. After the row-wise convolution, we can repeat the moving average operations on each column of the intermediate image. More information on our boxcar filtering algorithm can be found elsewhere.[20]

RESULTS AND DISCUSSION

We have achieved fast response times for the 1D real FFT and boxcar filtering, the two most computationally expensive algorithms. Our 1D 4096-point radix-2 complex FFT has an execution time of 371 ms for 2688 lines. In contrast, our 1D real FFT takes 215 ms. Thus, by performing the 1D real FFT, an execution time reduction of 42% was achieved. The processing speed of boxcar filtering is almost independent of kernel size, as shown in Figure 3. Increasing the kernel size from 7×7 to 201×201 results in

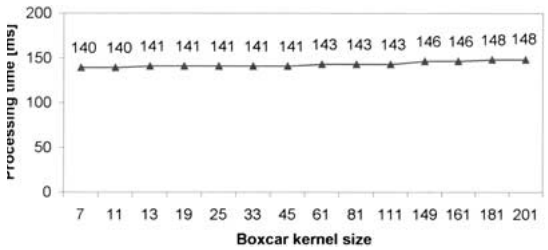

Fig 3. Performance of 2D boxcar filtering on a 2688 × 2688 12-bit image as a function of the boxcar kernel size.

Table 1. Execution time of enhancement tasks on a 2688 × 2688 12-bit image on a 450-MHz MAP processor.

Tasks	Time (ms)	Parameter
Grid suppression	559	
Adaptive unsharp masking	308	201 × 201 boxcar kernel
Rotate	38	90°
Window/level	25	12-bit to 8-bit
Image flip	36	
Lookup table	94	12-bit

only a 6% increase in processing time (from 140 to 148 ms), which is possible since each output after the initial moving average calculation of a row or column necessitates only a few arithmetic operations.

Table 1 shows the execution times attained for enhancement functions for a 12-bit 2688 × 2688 image. The grid suppression task is most computationally expensive at 559 ms, with the 1D real FFT, bandpass filtering, and the 1D real inverse FFT consuming the majority of cycles. Unsharp masking performance mainly comprises boxcar filtering and adaptively determining the weighting factor for each pixel. This adaptive filtering step, although capable of generating images of better quality, takes an additional 160 ms. While a direct comparison of our computing module with current CR and DR systems is difficult due to variations of the enhancement algorithms, our algorithms are representative of typical CR and DR enhancement methods. With the combined execution time of 465 ms for adaptive unsharp masking, window/level, image rotate, and LUT tasks, we have achieved at least an order of magnitude improvement compared with the response time of about 10 s in the DR systems.

There are currently three user-adjustable parameters to our image enhancement algorithms: the boxcar kernel size and the threshold for noise discrimination with adaptive unsharp masking as well as the LUT parameter. In addition, the user can control rotation and flip in the horizontal and vertical directions as well as the window width and level. The parameters of EVP, DRC, MUSICA, and other postprocessing algorithms could be mapped into our computing module or supported by a new module. For example, the input parameter *EVP Kernel Size* designates the size of a square boxcar kernel, which is typically set to 1/20th the size of the shorter of the width or height of the radiograph.[9] This equates to a 128 × 128 boxcar kernel size for a 2560 × 3072 image, which we can compute in 157 ms. With MUSICA, Vuylsteke and Schoeters[11] suggested the use of a LUT as a computationally inexpensive way to perform the enhancement operation. The *MUSI-contrast* parameter determines the amount of nonlinearity in a LUT, such that low-contrast details are amplified relative to high-contrast ones. For a 2688 × 2688 image, we measure the response time of 94 ms for a LUT operation.

Our computing module enables technologists or radiologists to view the output of a preset immediately. This allows them to interactively experiment with multiple presets, giving them the latitude to optimize the CR/DR image based on multiple factors, eg, examination type, patient size, dose, and pathology. For example, adaptive unsharp masking has an execution time of 308 ms for a 201 × 201 kernel (32 mm wide with a 0.16-mm pixel size). Thus, a radiologist can interactively view the result of different boxcar kernel sizes and noise thresholds in trying to improve the conspicuity of lung nodules and micronodule clusters in chest radiographs, where large boxcar kernels (25–70 mm) have been found superior.[21] At 25 ms, our window/level function can support real-time processing of CR/DR images in software without any hardware accelerator. This is important considering that window/level is the most common image enhancement function used by radiologists at PACS workstations.[22]

Postprocessing by technologists and image enhancement by radiologists are both critical links in the workflow chain, and the fast re-

sponse time of our enhancement algorithms will be very valuable. In Step 5 of Figure 1, the postprocessing performed by the technologist dictates when (1) the patient can be released and (2) the radiologist can review the radiograph on the PACS workstation. Thus, slow postprocessing by the technologist can prolong time to diagnosis and increase inconvenience to patients as well as tie up the examination room unnecessarily. In Step 7, fast execution of enhancement functions can decrease the diagnosis time of radiologists. While in this article we address two important steps of the workflow, improvements in the other steps can further enhance productivity. Progress in the accuracy of voice recognition software (Step 8), eg, can improve overall efficiency.

Our methods to expeditiously perform CR and DR image enhancement potentially have the greatest impact in a real-time radiology setting. With traditional radiology department setups, QA processing can create a bottleneck in the workflow based on the technologist, especially for CR if cassettes are not processed in a timely manner. On the other hand, the DR images are made available soon after X-ray exposure. Thus, after a quick check of the DR image by the technologist for positioning and orientation, with real-time radiology, the image and patient record can be automatically routed to the PACS workstation of a radiologist for immediate diagnosis. Thus, the technologist postprocessing step (Step 5) can now be removed from the workflow of Figure 1. While shifting the image postprocessing step to radiologists could represent additional work, the main advantage is that primary diagnosis can be made even faster, potentially before the patient leaves the examination room. The real-time delivery of the radiologist's interpretation ensures that the referring clinician views the images and report simultaneously, preventing situations where the clinician reviews the images and makes decisions without the radiologist's input, which can compromise patient care.[6] With real-time radiology, the radiologist receives raw images that have only been preprocessed and, coupled with our fast postprocessing, can apply personalized preset values or interactively process the images. The computing speed of modern processors will benefit radiologists via the quick enhancement of radiographs in making primary diagnosis.

Many commercially available CR/DR QA and PACS workstations currently use the processor in a personal computer, eg, Intel Pentium 4, to support image enhancement algorithms. For example, the Fuji QA Workstation 771 and GE Centricity Workstation both use Pentium-based processors. While image processing algorithms' execution times have been steadily decreasing due to improvements in processor architecture and increasing clock frequency, a further speedup by a factor of up to 10 could be realized by fully utilizing the capabilities of these processors, eg, MMX (multimedia extensions) and SSE (streaming data-level parallelism extensions).[15]

Mediaprocessors can be used for other computationally expensive tasks that can improve radiology productivity beyond CR/DR. Digital fluoroscopy and mammography are examples of other X-ray modalities in which mediaprocessors could speed up image enhancement. When executing our algorithms on digital mammography images, we can expect a response time proportional to the number of pixels, eg, with the 1900×2300 matrix size of the GE Senographe 2000D mammography system[23], we would expect a faster execution time than we report in Table 1. The same mediaprocessor used in a PACS workstation for image enhancement can be utilized to locally decompress an image retrieved from the PACS short-term archive before being displayed. We have developed a fast JPEG 2000 decoder on the MAP and obtained the execution time of 1.88 s for a 10:1 compressed, 12-bit, 2688×2688 DR image.[24] Computer-aided diagnosis (CAD) of CR/DR and other medical images is another example where mediaprocessors can be used to perform compute-intensive image processing and analysis operations, eg, segmentation, texture analysis, and classification.[25]

CONCLUSIONS

We have shown the feasibility of using a mediaprocessor as the main computing engine of CR/DR postprocessing algorithms. We efficiently mapped these algorithms to a single mediaprocessor and demonstrated the fast

computation of presets and/or interactive processing of full-resolution radiographic images. With the increasing capability and clock frequency of mediaprocessors, it will be possible to support new, more sophisticated algorithms and the increasing matrix size of CR/DR and/or other medical images.

We believe these modern processors, efficiently utilized in CR/DR devices and/or PACS workstations coupled with the reengineered workflow described in this article, will have a positive effect on productivity in the future. As radiology departments face the continual need to improve their productivity, the delivery of radiology service in real-time mode will become increasingly important. The use of these powerful processors will facilitate real-time radiology by minimizing the amount of time radiologists must spend in manipulating images, ensuring higher productivity and improved patient care in the future.

REFERENCES

1. Huang, HK: PACS: Basic Principles and Applications. New York: Wiley, 1999

2. Cao, X, Huang, HK: Current status and future advances of digital radiography and PACS. IEEE Eng Med Biol Mag 19:80-88, 2000

3. Freedman, MT, Artz, DS: Digital radiology. In: Siegel, EL, Kolodner, RM eds Filmless Radiology. Springer-Verlag, New York, 1999, pp 157-193

4. Reiner, BI, Siegel, EL: Technologists' productivity when using PACS: comparison of film-based versus filmless radiography. AJR Am J Roentgenol 179:33-37, 2002

5. Siegel, EL, Reiner, BI: Work flow redesign: the key to success when using PACS. Am J Roentgenol 178:563-566, 2002

6. Williamson, B: Real-time radiology. J Digit Imaging 12:48, 1999

7. http://www.agfa.com/healthcare/content/pdf/dr_thorax_en.pdf

8. http://www.kodak.com/global/en/health/productsByType/dr/dr5100.jhtml

9. Van Metter, R, Foos, D: Enhanced latitude for digital projection radiography. SPIE Proc 3658:468-483, 1999

10. Ishida M: Fuji computed radiography technical review, no. 1. Fuji Photo Film Co., Ltd., Tokyo, 1993

11. Vuylsteke, P, Schoeters, E: Multi-scale image contrast amplification (MUSICA) SPIE Proc 2167:551-560, 1994

12. Basoglu, C, Kim, D, Gove, RJ, et al: High-performance image computing with modern microprocessors. Int J Imaging Syst Technol 9:407-415, 1998

13. Basoglu, C, Lee, W, O'Donnell, J: The Equator MAP-CA TM DSP: an end-to-end broadband signal processor TM VLIW. IEEE Trans Circuits Syst Video Technol 12:646-659, 2002

14. Barski, LL, Wang, X: Characterization, detection and suppression of stationary grids in digital projection radiography imagery. SPIE Proc 3658:502-519, 1999

15. Grow, MS, Kim, Y: Evaluation of the Pentium 4 for imaging applications. SPIE Proc 4674:42-50, 2002

16. Lee, R: Accelerating multimedia with enhanced microprocessors. IEEE Micro 15:22-32, 1995

17. Porat, B: A Course in Digital Signal Processing Toronto: John Wiley & Sons, 1997

18. Mermer, C, Kim, D, Kim, Y: Efficient 2D FFT implementation on mediaprocessors. Parallel Computing 29:691-709, 2003

19. Smith, WW, Smith, JM: Handbook of Real-Time Fast Fourier Transforms. New York: IEEE Press, 1995

20. Bae, U, Shamdasani, V, Managuli, R, et al: Fast adaptive unsharp masking with programmable mediaprocessors. J Digit Imaging 16:230-239, 2003

21. Prokop, M, Schaefer, CM, Oestmann, JW, et al: Improved parameters for unsharp mask filtering of digital chest radiographs. Radiology 187:521-526, 1993

22. Horii, SC: Workstations. In: Dreyer, KJ, Mehta, A, Thrall, JH eds. PACS: A Guide to the Digital Revolution Springer-Verlag, New York, 2002, pp 191-235

23. http://www.gemedicalsystems.com/index.html

24. Agarwal, A, Rowberg, A, Kim, Y: Fast JPEG 2000 decoder and its use in medical imaging. IEEE Trans Information Technol Biomed 7:184-190, 2003

25. Van Ginneken, B, Ter Haar Romeny, BM, Viergever, MA: Computed-aided diagnosis in chest radiography: a survey. IEEE Trans Med Imaging 20:1228-1241, 2001